Application of perturbation theory to the stability analysis of realistic atmospheric flows

Mark D. Borges and Prashant D. Sardeshmukh

Cooperative Institute for Research in Environmental Sciences (CIRES)
University of Colorado, Boulder, CO 80309, USA

submitted *TELLUS*

October, 1996

ABSTRACT

The first order perturbation technique is investigated as a tool for measuring the dynamical significance of a change in a background flow. The focus is on estimating the change in the behavior of the linear disturbance operator as manifest in both the most rapidly growing normal modes as well as initial disturbances that amplify the most over finite time intervals (also known as singular vectors). The notion of a "small change" to a flow can be made more precise by requiring, for instance, that the changes in the flow's stability properties be accurately estimated using first-order perturbation theory. The perturbation theory updates of the eigen- and singular values are economical to compute, thus making this determination feasible even in a practical setting. Furthermore, a simple refinement of the basic algorithm enables an efficient updating of the singular vectors as well.

The technique is illustrated and tested for a flow having a substantial number of degrees of freedom, the global 250mb flow during northern hemisphere winter represented in a T31 spectral space. A complete generalized barotropic stability analysis of the observed long-term mean 250mb flow is performed first. Various alterations to the flow are then considered, and the perturbation technique is applied to estimate both the asymptotic and finite-time stability of the altered flows. Finally, the accuracy of these estimates is checked against the complete stability analyses of the altered flows, and an assessment is made of what a "small change" is for these cases.

It is shown that in many cases the perturbation theory estimates of the stability changes are accurate enough over a sufficiently large parameter range to be of practical use. Applications to the problems of: (1) anticipating variations in forecast skill associated with day-to-day variations in flow stability; and (2) anticipating the relevance and robustness of individual normal modes are discussed. The latter problem is closely related to the concept of pseudospectra, and perturbation theory can be used to estimate their coarse details.

1 Introduction

The value of an atmospheric model is ultimately determined by its ability to serve as a proxy for the real atmosphere. Indeed, the underlying premise of many hypothetical model experiments is that the results are applicable to some aspect of the real atmosphere *e.g.*, to some spatial or temporal region of the real atmosphere's statistics. Accordingly, model validation is one of the more important components of any modeling effort, yet it is often the least objectively determined measure of a model.

The simplest model validation is done using gross statistical measures. For instance, an assessment of a model's performance in simulating present-day climatology is typically made by comparing mean flow or variance maps produced by the model with their observed counterparts. This provides the investigator with only a crude indication, however, of whether the model simulation is a "reasonable" facsimile of the observed climatology. A more precise determination of the significance of the differences requires the use of some objective method. If one is interested in the differences in the mean flow itself, one could use standard statistical techniques to determine if the differences are statistically significant. But perhaps one is also interested in how the change in the mean flow alters the evolution of perturbations to that flow. In other words, does a perturbation evolve in a significantly different manner on the model climatology compared to the observed climatology? In this latter case, we are interested in whether the differences are *dynamically* significant.

The distinction between statistical and dynamical significance can be clarified by considering the following hypothetical experiment. Suppose the difference in the basic state is such that it translates to a simple multiplicative change in the linear operator, $\mathbf{L}$, that governs the linearized dynamics of the system. It can be shown for this simplified case that the eigenfunctions of $\mathbf{L}$ do not change; only the eigenvalues scale by the multiplicative factor. Since the linear eigenfunction evolution would remain unchanged, from a dynamical perspective the basic state modification would be irrelevant. Therefore, the difference could conceivably be made statistically significant and yet remain dynamically insignificant.

Although somewhat contrived, the above scenario illustrates the difficulty of assessing the dynamical impact of a change in the basic state. A large change could conceivably have no effect on the linearized dynamics, while conversely a relatively small change in a particular location of the basic state could lead to significantly different perturbation evolution. One of the main points of this paper is to address this difficulty and suggest a dynamical definition

for basic flow changes: we consider a basic state modification to be large from a dynamical perspective if it substantially alters the stability properties of the associated $\mathbf{L}$ operator.

An efficient, objective method to determine the dynamical significance of basic state differences is clearly needed. This paper will demonstrate that the first-order perturbation technique (*see, for example,* Wilkinson, 1965; North *et al.*, 1982; Penland and Sardeshmukh, 1995) provides such a tool. The technique can also be used to address two other common questions arising in a stability analysis:

(i) how reliable or sensitive are the stability analysis results?
(ii) how meaningful are the individual modes?

The computational effort required to answer *(i)* and *(ii)* in an applied setting will also be discussed.

To fix ideas, consider the archetypal linear system,

$$\frac{d}{dt}\underset{\sim}{\mathbf{x}} = \mathbf{L}\underset{\sim}{\mathbf{x}} + \underset{\sim}{\mathbf{f}} \tag{1}$$

where $\underset{\sim}{\mathbf{x}} = \underset{\sim}{\mathbf{x}}(t)$ is the state vector of the system, $\mathbf{L}$ is a matrix representing the relevant discretized linear operator and $\underset{\sim}{\mathbf{f}} = \underset{\sim}{\mathbf{f}}(t)$ is a time-dependent forcing function vector. Our interest is in how the $\mathbf{L}$ operator behaves as perturbations are made to it. In this sense, we use the term "behavior" to refer to the sensitivity or robustness of the results of a stability analysis of $\mathbf{L}$. This sensitivity may arise not only from the nature of the problem, but also due to inherent uncertainties in the linearization. For example, because of errors in the formulation of a model, or simply not knowing the correct basic state to linearize about, the precise definition of the operator $\mathbf{L}$ may not be well known. Say

$$\hat{\mathbf{L}} = \mathbf{L} + \delta\mathbf{L} \tag{2}$$

where $\hat{\mathbf{L}}$ is the estimate of the true operator $\mathbf{L}$. Knowing how the change $\delta\mathbf{L}$ affects the properties and behavior of the true $\mathbf{L}$ is at least as important as the results of the analysis on $\hat{\mathbf{L}}$ itself.

Traditionally, much emphasis has been placed on the asymptotic stability properties of (1) as governing its behavior. For instance, exponential (free) solutions of the linear equation (1) are sought, resulting in an eigenvalue problem,

$$\mathbf{L}\underset{\sim}{\mathbf{u}}_j = \lambda_j \underset{\sim}{\mathbf{u}}_j \tag{3}$$

for the complex eigenvector $\underset{\sim}{\mathbf{u}}_j$ and eigenvalue $\lambda_j = (\sigma_j + i\omega_j)$. If the operator $\mathbf{L}$ is self-adjoint the eigenvectors and eigenvalues are pure real, and the eigenvectors can be normalized to form an orthonormal spectral basis. In this

situation eigenspectra do, in fact, provide valuable information about the time evolution of perturbations governed by (1). Furthermore, it can be shown that a change $\delta\mathbf{L}$ of a given size can be used to determine a bound on the change in the eigenvalues, without the need to repeat the complete eigenvalue calculation.

Unfortunately, for many meteorological and hydrodynamical applications the $\mathbf{L}$ operator in (1) is *not* self-adjoint. The associated eigenvectors do not form an orthogonal spectral basis by themselves. For these non-normal operators eigenanalysis may fail to give an accurate indication of the behavior of $\mathbf{L}$, and as discussed later, may even be deceptive. Also, to our knowledge there is no simple method to determine a bound on the change in the eigenvalues given a change $\delta\mathbf{L}$ in a non-normal $\mathbf{L}$.

As alluded to previously, the case when $\mathbf{L}$ is self-adjoint simplifies the assessment of the significance of the basic state change. The size of a given change in the matrix operator can be used to find a bound on the change possible in a particular stability parameter (*e.g.,* the leading eigenvalue). Discussion of various aspects may be found in the classic work of Wilkinson (1965), and in a meteorological context in North *et al.* (1982).

More recently, Penland and Sardeshmukh (1995, hereafter PS95) have discussed the extension of first-order perturbation theory to non-normal systems and how it may be used to determine the modifications to σ_j, ω_j, and $\underset{\sim}{\mathbf{u}}_j$ caused by a small change in the basic state.

To help understand question *(ii)*, consider the forced problem (1) with $\underset{\sim}{\mathbf{f}}(t) = \exp(\mu t)\underset{\sim}{\mathbf{f}}_0$. The general solution may be written,

$$\underset{\sim}{\mathbf{x}}(t) = \exp(t\mathbf{L})[\underset{\sim}{\mathbf{x}}_0 - (\mu\mathbf{I}-\mathbf{L})^{-1}\underset{\sim}{\mathbf{f}}_0] + \exp(\mu t)(\mu\mathbf{I}-\mathbf{L})^{-1}\underset{\sim}{\mathbf{f}}_0 \tag{4}$$

where $\underset{\sim}{\mathbf{x}}(0) \equiv \underset{\sim}{\mathbf{x}}_0$, $\mathbf{I}$ is the identity matrix and $\mu \in C$. As is well known, the solution, (4), will exhibit resonant behavior if μ happens to be an eigenvalue of $\mathbf{L}$ (*i.e.,* $\mu \in \Lambda(\mathbf{L})$)**.** But it will also exhibit resonant-like behavior if the quantity $(\mu\mathbf{I}-\mathbf{L})^{-1}$ is large in some sense. Hence values of $\mu \in C$ for which $\|(\mu\mathbf{I}-\mathbf{L})^{-1}\|$ is large, where $\|\ \|$ is any chosen norm, are also important. In other words, it is equally meaningful to ask how quickly $\|(\mu\mathbf{I}-\mathbf{L})^{-1}\|$ decays away from an eigenvalue. If the decay is slow, the emphasis on the eigenmodes of the problem may be unduly exaggerated. To address this issue, one can appeal to the notion of *pseudospectra* or *spectral portraits* (see, for example, Trefethen

et al. 1993). Alternatively, if a sufficiently "small" change is considered, first-order perturbation theory can be used for substantial computational savings.

Finally, it should be noted that the implementation of the first-order perturbation theory need not be computationally inefficient even for problems that require updating eigenvectors. While it is true that the updating of eigenvectors via perturbation theory is an $O(N^3)$ operation (where N is the rank of the matrix), some operators permit a simplification to the basic algorithm that greatly diminishes the number of floating point operations required, and without significant loss of accuracy. This result makes feasible the possibility of using perturbation theory to determine certain stability properties on a daily basis in an inexpensive and efficient manner.

The paper is organized as follows. Section 2 reviews the formulation of the first-order perturbation theory formulation, drawing on the previous work of North, *et. al* (1982) and PS95. In the following section, results for two hypothetical cases are presented, illustrating how well the theory works in a system with a large number of degrees of freedom. Section 4 addresses computational efficiency considerations, and introduces a refinement of the basic algorithm that works well for some problems. In section 5 we discuss the concept of pseudospectra and how it relates to this work. A summary and concluding remarks are presented in section 6.

2 Governing Equations

In this section the perturbation technique is illustrated for a simple system with a substantial number of degrees of freedom — the barotropic vorticity equation. A more complete description of the linearization and notation may be found in Borges and Sardeshmukh (1995). For convenience, a brief version follows.

The barotropic vorticity equation may be written

$$\frac{\partial \zeta}{\partial t} + \nabla \bullet (\underset{\sim}{\mathbf{V}}_\psi \, \zeta) \; = \; F \tag{5}$$

where ζ is the absolute vorticity and $\underset{\sim}{\mathbf{V}}_\psi$ is the horizontally nondivergent wind. Separating variables into a time-independent basic state and a time-dependent perturbation, equation (1) becomes,

$$\frac{\partial}{\partial t}\zeta' + \nabla \bullet \bar{\underset{\sim}{\mathbf{V}}}_\psi \zeta' + \nabla \bullet \underset{\sim}{\mathbf{V}}'_\psi \; \bar{\zeta} = \; F - \nabla \bullet \bar{\underset{\sim}{\mathbf{V}}}_\psi \bar{\zeta} \equiv F' \equiv 0 \tag{6}$$

where quadratic terms in perturbation amplitude are neglected, and we have made the "SWB" simplifying assumption

in setting the rhs of (6) to zero as described in Borges and Sardeshmukh (1995). Expanding in a truncated series of spherical harmonics, this may be cast into an ordinary differential equation,

$$\frac{d}{dt}\underset{\sim}{\zeta}' = \mathbf{L}\,\underset{\sim}{\zeta}' \tag{7}$$

The operator, $\mathbf{L}$, is determined at T31 truncation giving a real matrix of rank N = 1023.

Exponential solutions to the linear equation (7) are sought, resulting in the eigenvalue problem,

$$\mathbf{L}\underset{\sim}{\mathbf{u}}_j = \lambda_j \underset{\sim}{\mathbf{u}}_j \tag{8}$$

for the complex eigenvector $\underset{\sim}{\mathbf{u}}_j$ and eigenvalue $\lambda_j = (\sigma_j + i\omega_j)$. Since $\mathbf{L}$ is in general not self-adjoint, we consider the eigenvalue problem for the adjoint,

$$\mathbf{L}^{\dagger}\underset{\sim}{\mathbf{v}}_j = \lambda^*_j \underset{\sim}{\mathbf{v}}_j \tag{9}$$

where the asterisk denotes a complex conjugate and the adjoint operator $\mathbf{L}^{\dagger}$ is defined by $\langle \underset{\sim}{\mathbf{v}}_j, \mathbf{L}\underset{\sim}{\mathbf{u}}_j \rangle = \langle \mathbf{L}^{\dagger}\underset{\sim}{\mathbf{v}}_j, \underset{\sim}{\mathbf{u}}_j \rangle$ for a suitable scalar product $\langle \bullet,\bullet \rangle$. We will use a scalar product proportional to rotational kinetic energy,

$$\langle f, g \rangle = \int_{-\pi/2}^{\pi/2} \cos\theta d\theta \left(\frac{1}{2\pi} \int_0^{2\pi} \nabla f \bullet \nabla g d\lambda \right) \tag{10}$$

The importance of the adjoint vectors $\underset{\sim}{\mathbf{v}}_j$ is that they form a bi-orthogonal set with the eigenvectors $\underset{\sim}{\mathbf{u}}_j$, *i.e.*, $\langle \underset{\sim}{\mathbf{v}}_j, \underset{\sim}{\mathbf{u}}_k \rangle = \delta_{jk}$. The complex eigenvector $\underset{\sim}{\mathbf{u}}_j$, may be expressed in geographical space in the form,

$$\underset{\sim}{\mathbf{u}}_j = [\underset{\sim}{\mathbf{a}}_j(\lambda, \theta)\cos\omega_j t + \underset{\sim}{\mathbf{b}}_j(\lambda, \theta)\sin\omega_j t]e^{\sigma_j t} \tag{11}$$

where $\underset{\sim}{\mathbf{a}}_j$ and $\underset{\sim}{\mathbf{b}}_j$ are normalized functions of longitude and latitude that are orthogonal to each other in space, *i.e.*,

$$\langle \underset{\sim}{\mathbf{a}}_j, \underset{\sim}{\mathbf{b}}_j \rangle \equiv 0\,, \qquad \langle \underset{\sim}{\mathbf{a}}_j, \underset{\sim}{\mathbf{a}}_j \rangle = 1 \geq \langle \underset{\sim}{\mathbf{b}}_j, \underset{\sim}{\mathbf{b}}_j \rangle\,. \tag{12}$$

2.1 Perturbation Theory

Perturbation theory permits one to determine how σ_j, ω_j, $\mathbf{a}_j$ and $\mathbf{b}_j$ are modified as the basic state is changed. The derivation below closely follows that of PS95.

Setting $\mathbf{L} = \mathbf{L} + \delta\mathbf{L}$, $\lambda_j = \lambda_j + \delta\lambda_j$, $\underset{\sim}{\mathbf{u}}_j = \underset{\sim}{\mathbf{u}}_j + \delta\underset{\sim}{\mathbf{u}}_j$ in (8) and neglecting quadratic perturbation quantities one

obtains,

$$(\mathbf{L} - \lambda_j \mathbf{I})\delta\underset{\sim}{\mathbf{u}}_j \;=\; (\delta\lambda_j \mathbf{I} - \delta\mathbf{L})\underset{\sim}{\mathbf{u}}_j \tag{13}$$

Expanding $\delta\underset{\sim}{\mathbf{u}}_j$ in the (presumably complete) set of eigenfunctions, $\{\underset{\sim}{\mathbf{u}}_k\}$ provided by the numerical discretization gives

$$\delta\underset{\sim}{\mathbf{u}}_j \;=\; \sum_{\substack{k \\ k \neq j}} c_{jk}\underset{\sim}{\mathbf{u}}_k \tag{14}$$

where the summation is over all eigenmodes except the j^{th} eigenmode. This exclusion is permissible because $(1 + c_{jj})\underset{\sim}{\mathbf{u}}_j \;=\; \hat{\underset{\sim}{\mathbf{u}}}_j$ is still an eigenvector of $\mathbf{L}$, (the normalization of the eigenvectors is irrelevant up to this point) and $\hat{\underset{\sim}{\mathbf{u}}}_j$ can be substituted for $\underset{\sim}{\mathbf{u}}_j$ anywhere above without loss of generality. Substituting (14) in (13) yields,

$$(\mathbf{L} - \lambda_j \mathbf{I}) \sum_{\substack{k \\ k \neq j}} c_{kj}\underset{\sim}{\mathbf{u}}_k \;=\; (\delta\lambda_j \mathbf{I} - \delta\mathbf{L})\underset{\sim}{\mathbf{u}}_j \tag{15}$$

Forming the inner product of (15) with $\underset{\sim}{\mathbf{v}}_n$ gives

$$\langle \underset{\sim}{\mathbf{v}}_n, (\mathbf{L} - \lambda_j \mathbf{I}) \sum_{\substack{k \\ k \neq j}} c_{kj}\underset{\sim}{\mathbf{u}}_k \rangle \;=\; \langle \underset{\sim}{\mathbf{v}}_n, (\delta\lambda_j \mathbf{I} - \delta\mathbf{L})\underset{\sim}{\mathbf{u}}_j \rangle \tag{16}$$

For $n = j$, (16) reduces to

$$0 \;=\; \langle \underset{\sim}{\mathbf{v}}_j, (\delta\lambda_j \mathbf{I} - \delta\mathbf{L})\underset{\sim}{\mathbf{u}}_j \rangle \tag{17}$$

since the eigenvectors and their adjoints can be constructed as a bi-orthonormal set and $j \neq k$ in the summation appearing in the lhs of (16). Alternatively one could appeal to the Fredholm alternative, as in PS95, to obtain the same result without explicit use of the expansion (14).

Therefore, from (17) the change in the eigenvalue is,

$$\delta\lambda_j \;=\; \frac{\langle \underset{\sim}{\mathbf{v}}_j, \delta\mathbf{L}\underset{\sim}{\mathbf{u}}_j \rangle}{\langle \underset{\sim}{\mathbf{v}}_j, \underset{\sim}{\mathbf{u}}_j \rangle} \;=\; \langle \underset{\sim}{\mathbf{v}}_j, \delta\mathbf{L}\underset{\sim}{\mathbf{u}}_j \rangle \tag{18}$$

where the chosen normalization satisfies $\langle \underset{\sim}{\mathbf{v}}_j, \underset{\sim}{\mathbf{u}}_j \rangle \equiv 1$ and a perturbation in the basic flow, say $\delta\bar{\zeta}$, directly translates to a change $\delta\mathbf{L}$ in the linear matrix operator. The change in a particular eigenvector may be determined by solving for the expansion coefficients c_{kj} in (16). Substituting (18) into (16), using the bi-orthonormal property and simplifying

one obtains,

$$(\lambda_j - \lambda_n) c_{nj} = \langle \mathbf{v}_n, \delta \mathbf{L} \mathbf{u}_j \rangle, \quad n \neq j \tag{19}$$

The c_{jj} coefficients are implicitly determined upon normalization of the updated eigenvector $\mathbf{u}_j = \mathbf{u}_j + \delta \mathbf{u}_j$.

3 Results

In this section, a set of basic states is systematically varied and analyzed using eigenanalysis and singular value decomposition (SVD). The perturbation theory estimate of the change in the stability properties is then compared with the change determined by solving the full eigenvalue or SVD problem using standard matrix algorithms. The basic state changes are characterized by a single parameter. The parameter range for which the perturbation theory fails can be used to indicate when the corresponding change in the basic state can be considered dynamically significant.

The performance of the perturbation technique in capturing structural changes is assessed using both a global and local measure. The global measure is the "maximum pattern correlation", defined as in Borges and Sardeshmukh (1995, page 3787). The local measure is a normalized maximum local error in the expansion coefficient vector of the eigenvector, defined as

$$\| \hat{q} - q \|_\infty / \| q \|_\infty \tag{20}$$

where

$$\begin{aligned} u &= \mathbf{U} q \\ u^{(est)} &= \mathbf{U} \hat{q} \end{aligned} \tag{21}$$

and $\| \ \|_\infty$ is the infinity norm ($\|\mathbf{x}\|_\infty = \max_i |x_i|$). Here $\mathbf{U}$ is the eigenbasis matrix of the original matrix $\mathbf{L}$, u is the spectral coefficient vector of the actual updated eigenmode, and $u^{(est)}$ is the perturbation estimate of it. Note that the difference in the numerator of (20) is obtained after rotating $u^{(est)}$ by ϕ_{max} where the phase ϕ_{max} is determined by the maximum pattern correlation between u and $u^{(est)}$ (see Borges and Sardeshmukh, 1995).

3.1 Data source

The two data sources used for this investigation are the initialized analyses of horizontal wind obtained from the

European Centre for Medium Range Weather Forecasts (*ECMWF*) and the National Meteorological Center (*NMC*). The *ECMWF* dataset consists of 13 years (1979 through 1991) of global winter (DJF) data at 250mb. The *NMC* DJF dataset is based on 27 years (1965 through 1991) of Northern Hemisphere 250mb data and 14 years (1978 through 1991) of Southern Hemisphere data. The vector wind components were converted to scalar fields of vorticity and divergence and expanded in a T31 spherical harmonic basis, as in Borges and Sardeshmukh (1995). The climatological wintertime rotational flows used subsequently are defined as the time-average vorticity.

Figures 1a and 1b show the streamfunction and isotach representation of the *NMC* and *ECMWF* climatologies, respectively[1]. The experiments discussed in the following section evaluate the success of the perturbation method in estimating the eigenvectors and singular vectors of the *ECMWF* given those of the *NMC* climatology. The degree of success can be used to assess whether the differences, which appear minor to the eye, are in fact dynamically insignificant. In Section 3.3 the zonally asymmetric component of the *NMC* climatology is systematically varied, and the perturbation method is used to estimate the changes to the stability properties of these altered background flows. Figures 1c and 1d show two examples for the cases when the climatological nonzonal component is decreased and increased by 20%, respectively.

3.2 *NMC* and *ECMWF* blending experiment

The first set of experiments was designed to determine the significance of the differences between the *NMC* and *ECWMF* analyzed DJF climatologies. The global mean rotational kinetic energy of the difference between the *NMC* and *ECMWF* climatologies is ~1 m^2s^{-2}. The energy of the difference field is nearly equally partitioned between the hemispheres; 58% of the total energy (0.6 m^2s^{-2}) is contained in northern hemisphere latitudes.

To determine if these differences are dynamically significant, a blended basic flow was defined as

$$\bar{\zeta}_\gamma = (1-\gamma)\bar{\zeta}_{NMC} + \gamma\bar{\zeta}_{ECMWF} \tag{22}$$

and the parameter γ varied between 0.05 and 0.95 in discrete steps.

3.2.1 Eigenanalysis

Figure 2 shows the fractional change in the growth rate σ_1 and frequency ω_1 of the leading eigenvalue as a function of

1. A difference map of these climatologies is also shown in Borges and Sardeshmukh (1995).

γ. The fractional change in the growth rate (Fig. 2a) obtained by repeating the full matrix calculation for each basic state (hatched bars) is captured quite well by the perturbation theory estimate (solid bars). Conversely, the full matrix calculation indicates that the fractional change in the frequency (Fig. 2b) undergoes a sign reversal, which is not captured by the perturbation theory. However, the magnitude of the change is fairly small, an order of magnitude less than the change in the growth rate.

Figure 3 summarizes the magnitude of the structural changes in the leading eigenvector. The fractional error in the perturbation estimate (Fig. 3a) is largest for the **a**-phase (the most energetic phase) of the eigenmode, and approaches 0.45 near the ECWMF basic state at γ=0.95. The maximum error in the **b**-phase of the eigenmode is less, approaching 0.25 at γ=0.95. Figure 3b shows the pattern similarity (at the phase of the maximum pattern correlation) between the perturbation theory estimate of the structure and that obtained from the full matrix problem. For the range of γ considered, the maximum pattern correlation always exceeds 0.9.

According to Fig. 3b, the perturbation theory works well, whereas according to Fig. 3a it works less well. It should be borne in mind, however, that Fig. 3a emphasizes the maximum local differences in the eigenmode structures. We conclude that perturbation theory succeeds in accurately estimating the overall change in the leading normal mode's eigenstructure. Therefore, from the dynamical perspective of the leading eigenmode, the differences between the *NMC* and *ECMWF* basic states should not be considered significant.

3.2.2 Singular value decomposition

The above analysis was repeated for the singular value decomposition of $\mathbf{G} = \exp(\tau\mathbf{L})$, for τ = 3 days. The singular value decomposition is equivalent to eigenanalysis of the self-adjoint matrices $\exp(\tau\mathbf{L}^\dagger)\exp(\tau\mathbf{L})$ and $\exp(\tau\mathbf{L})\exp(\tau\mathbf{L}^\dagger)$. We obtain the complete singular value decomposition using a standard matrix technique available via the NetLib[2] software repository.

Figure 4 shows the fractional change in the singular values as a function of γ. The perturbation estimates decrease the singular value by about 20% as the *ECMWF* basic state is approached, nearly the same as that given by the full matrix computation of the singular values (~17%). The perturbation estimate systematically overestimates the actual change for all γ.

2. The NetLib repository is accessible via the World-Wide Web through the URL, "*http://www.netlib.org*".

The structural changes are summarized in Fig. 5, in the same format as in Section 3.2.1. The normalized error (Fig. 5a) is less than that for the leading normal mode in Fig. 3a. At γ=0.95, the error is about 19%, compared to ~45% for the **a**-phase leading normal mode. The pattern correlation between the estimated leading singular vector and the actual singular vector exceeds 0.96 for all γ, again better than that in Fig. 3b. For this self-adjoint system, the perturbation estimates are more accurate compared to the non-self-adjoint system.

Again, perturbation theory succeeds in accurately estimating the change in the leading singular value and singular structure. As such, the differences between the *NMC* and *ECMWF* basic states, from the dynamical perspective of the leading singular vector, are not significant.

3.3 Zonally asymmetric experiments

The second set of experiments was designed to examine the significance of the nonzonal component of the basic state. A set of basic states was constructed by systematically increasing the magnitude of the nonzonal component while keeping the zonal mean fixed, *i.e.*,

$$\bar{\zeta}_\alpha = [\bar{\zeta}_{NMC}] + \alpha\bar{\zeta}^*_{NMC} \tag{23}$$

where square brackets denote the zonal mean, the asterisk the departure from the zonal mean, and the parameter α is varied between 0.6 and 1.6. For reference, the global mean rotational kinetic energy of the difference of the *NMC* and *ECMWF* basic states (1 m^2s^{-2}), discussed in Section 3.2, is equivalent to an α-value of about 1 ± 0.2 . The energy of this difference field ($\bar{\zeta}_{\alpha=1.2} - \bar{\zeta}_{NMC}$), however, is not as equally partitioned between the hemispheres. For the $\bar{\zeta}_{1.2}$ basic state (Fig. 1d), more than 80% of the global mean rotational kinetic energy (0.87 m^2s^{-2}) is contained in the Northern Hemisphere. Furthermore, the local wind maxima in the ($\bar{\zeta}_{\alpha=1.2} - \bar{\zeta}_{NMC}$) difference field (not shown) are nearly double those of the *ECMWF* - *NMC* difference field.

3.3.1 Eigenanalysis

Figure 6 shows the fractional change in the growth rate and frequency of the leading normal mode as a function of α. The actual change in the growth rate (hatched bars, Fig. 6a) is asymmetric about α=1.0 (*i.e.*, the change for α = 1.3 is greater than that for α = 0.7) and levels off at about α = 1.4. The first order perturbation theory is unable to account for these features. The perturbation estimates are reasonably accurate for $0.7 \le \alpha \le 1.3$, and are slightly better for $\alpha >$

1 due to the asymmetry noted previously. For $\alpha > 1.5$, the perturbation theory predicts an $O(1)$ change, while the actual change is about 0.6.

The changes in the leading eigenvalue's frequency are shown in Fig. 6b. The actual changes in the frequency (hatched bars) are again asymmetric about $\alpha = 1.0$, but now in the opposite sense. Consequently, the perturbation theory result for $\alpha = 0.7$ is better than that for $\alpha = 1.3$. In contrast to the growth rate, the fractional changes in frequency remain below 0.5 for α in the range $0.7 \le \alpha \le 1.6$.

The perturbation estimates of the leading eigenmode structure (Fig. 7) are reasonably accurate for $0.7 \le \alpha \le 1.4$ The maximum pattern correlation (Fig. 7b) for this range of α exceeds 0.8. The normalized change in the expansion coefficient vector for the **a**-phase of the leading eigenmode (Fig. 7a, solid bars) is at 0.5 or below for $0.7 \le \alpha \le 1.5$. The degradation in the maximum pattern correlation is primarily due to the larger error in the less energetic **b**-phase of the leading eigenvector, particularly for $\alpha > 1.4$.

3.3.2 Singular value decomposition

The fractional change in the leading singular values for α in the range $0.7 \le \alpha \le 1.6$ is shown in Fig. 8. As in the γ-experiment, the perturbation estimates of the singular values are better than the leading eigenvalue over the entire range shown. For $\alpha \approx 1.4$ and above the perturbation theory predicts an $O(1)$ change to the singular value, a change that is also evident in the singular value obtained from the full problem.

The structural changes to the leading singular vector as a function of α are summarized in Fig. 9. The maximum pattern correlation exceeds 0.8 for α in the range $0.8 \le \alpha \le 1.3$ (Fig. 9b). The normalized error in the expansion coefficient vector (Fig. 9a) varies from 0.35 for $\alpha = 0.8$ to 0.5 for $\alpha = 1.3$.

4 Computational considerations

Perturbation theory updates of the eigenvalues of $\mathbf{L}$ are obtained using (18), while updates of the eigenvectors are obtained using (14) and (19). The eigenvalue updates are computationally inexpensive; for instance, a straightforward implementation of the algorithm is about a factor of two faster than the standard SVD solvers available through the NetLib software repository. At first sight, the eigenvector updates are not inexpensive. In general, one must perform the sum in (14) over the entire set of eigenfunctions. This becomes an $O(N^3)$ operation, which can be computationally

expensive. However, based on inspection of maps of the expansion coefficient amplitude, a simplification of (14) is possible that works well for updating the singular vectors. For example, Fig. 10 shows the average amplitude along each diagonal of the right singular vector expansion coefficient matrix as a function of the diagonal index, *l*. The notation is such that the average along the main diagonal is labelled by the $l = 0$ index, and positive or negative diagonal indices correspond to averages along diagonals in the upper right or lower left triangular portion of the matrix, respectively. The expansion coefficient matrix is sparse, with non-negligible amplitude in only a few tens of elements about the main diagonal. Figure 10 suggests that limiting the summation in (14) to a band, *J*, about the original singular vectors, *i.e.*,

$$\delta\underset{\sim}{\mathbf{u}}_j^{(J)} = \sum_{\substack{k = max(1,\, j-J) \\ k \neq j}}^{min(N,\, j+J)} c_{jk}\underset{\sim}{\mathbf{u}}_k \tag{24}$$

will give a very good approximation of the actual perturbation theory estimate of $\delta\underset{\sim}{\mathbf{u}}_j$. Accordingly, we have recomputed the ten leading $\delta\underset{\sim}{\mathbf{u}}_j's$ using this simplification with *J*=10 and found that the correlation between $\delta\underset{\sim}{\mathbf{u}}_j^{(10)}$ and $\delta\underset{\sim}{\mathbf{u}}_j^{(1022)}$ (the original perturbation estimate obtained using *J*=1022) exceeds 0.98 in all cases. Given the substantial computational savings and minimal loss of accuracy, this simplification may prove useful in obtaining preliminary estimates of the change for more complex operators with a large number of degrees of freedom.

5 Discussion

We have shown how first-order perturbation theory may be used to assess the dynamical significance of a change to the basic state, by examining to what extent the change in its stability properties can be captured by perturbation theory. If the change is not accurately captured by the perturbation theory, one may regard the basic state change as dynamically significant.

Perturbation theory can also be useful in generating "poor man's" *pseudospectra*. The pseudospectrum of an operator $\mathbf{L}$ extends the concept of its eigenspectrum, $\Lambda(\mathbf{L})$, defined as the set of all eigenvalues λ in the complex plane *C*. In an analogous manner, the ε-pseudospectra $\Lambda_\varepsilon(\mathbf{L})$ of $\mathbf{L}$ are defined as,

$$\Lambda_\varepsilon(\mathbf{L}) = \{z \in C:\ \left\|(z\mathbf{I} - \mathbf{L})^{-1}\right\| \geq \varepsilon^{-1}\} \tag{25}$$

with the convention that $\left\|(z\mathbf{I} - \mathbf{L})^{-1}\right\| = \infty$ if $z \in \Lambda(\mathbf{L})$ as in Trefethen (1992). In a forced problem like (1) with

solution (4), apparent resonant-like behavior of the system may be observed if $\|(\mu\mathbf{I}-\mathbf{L})^{-1}\|$ is large, even though it may be far from resonance in the traditional sense. The pseudospectrum provides valuable information in such systems by delimiting the region of particular values of μ in the complex plane where a large response to forcing should be anticipated. Thus, eigenanalysis alone of non-normal operators may fail to give an accurate indication of the behavior of the $\mathbf{L}$ operator, and worse, can be misleading to the investigator.

As noted by Trefethen (1992), it can be shown that an equivalent definition of $\Lambda_\varepsilon(\mathbf{L})$ is,

$$\Lambda_\varepsilon(\mathbf{L}) = \{z \in C:\ z \in \Lambda(\mathbf{L}+\delta\mathbf{L}) \text{ with } \|\delta\mathbf{L}\| \le \varepsilon\} \tag{26}$$

Viewed from this perspective, the pseudospectrum can be used to visually reveal which eigenvalues are sensitive to changes in $\mathbf{L}$. For self-adjoint systems, $\Lambda_\varepsilon(\mathbf{L})$ is described by the union of circles of radius ε about each eigenvalue. In non-normal systems, however, the shape of $\Lambda_\varepsilon(\mathbf{L})$ can differ substantially from this simple form, implying that a relatively small change in the basic state could lead to a much larger change in the eigenvalue under consideration. Thus, in an initial value problem as well as the forced problem the investigator must be wary of overemphasizing the importance of a specific normal mode.

The pseudospectrum is a powerful analysis tool, but in practice the computation of Λ_ε can be prohibitive. In principle, the determination of the resolvent norm, $\|(z\mathbf{I}-\mathbf{L})^{-1}\|$, must be repeated for each point z in the complex domain of interest. However, from (26) the relation to the perturbation theory discussed here is readily apparent. When the operator is not degenerate, information akin to pseudospectra may thus be obtained by appealing to perturbation theory. Unlike the full Λ_ε computation, the approximate Λ_ε obtained via the perturbation technique may be computationally feasible for a wider range of problems.

6 Summary and Concluding Remarks

Perturbation theory is a useful tool for objectively measuring the dynamical significance of a change in a basic flow. Our definition of dynamical significance is based upon the change in the behavior of the linear disturbance operator as manifest in the most rapidly growing normal modes and singular vectors. The notion of a “small” or dynamically insignificant change to a flow can be made more precise by requiring, for instance, that the changes in the flow’s stability properties be accurately estimated using first-order perturbation theory. The perturbation theory updates of the

eigen- and singular values are economical to compute, thus making this determination feasible even in a practical setting. Our principal findings may be summarized as follows.

- Although the perturbation theory performed well in the experiments where the *NMC* climatology was changed to the *ECMWF* climatology (*i.e.*, the blending experiments), it performed even better in the experiments with the nonzonal component of the *NMC* climatological flow changed by 20%, where the change was more noticeable (*i.e.*, the α-experiments with $\alpha = 1 \pm 0.2$). This illustrates the difficulty in trying to assess the dynamical significance of basic state differences by visual inspection alone, a task that is better suited to the perturbation method.
- The perturbation theory provides useful estimates of the stability parameters for changes to the nonzonal component (with fixed zonal mean) of about 20%. For this type of basic state modification, changes greater than about 30% should be considered dynamically significant for perturbations.
- For some non-degenerate problems, “poor man’s pseudospectra” may be obtained using perturbation theory.

The success of perturbation theory as a way to accurately estimate stability properties in a practical setting obviously depends upon the size of the basic state changes considered. Our experience suggests that a basic state change of $\sim 1\ m^2 s^{-2}$ in the globally averaged rotational kinetic energy is small enough that stability parameter modifications can be accurately approximated using the perturbation method. If the daily departures from a climatological basic state (say a seasonal or monthly mean) lie within this limit, the perturbation method can be used as an efficient and powerful method of estimating daily fluctuations of stability properties. This possible use of the perturbation technique is the subject of ongoing research, and could prove beneficial to operational centers that compute stability properties on a routine basis. For example, Buizza and Palmer (1995) recently investigated the daily fluctuations of atmospheric singular values obtained from the ECMWF Integrated Forecasting System using a Lanczos method. Since singular vectors bound the maximum possible growth of all perturbations, this approach of varying the basic state and computing the changed singular vector structure can lead to increased understanding of the model’s potential for error growth and how that error growth potential varies with time due to the changing background flow.

The perturbation method is also a useful tool for estimating the structural gradient of stability parameters with respect to the background flow. The structural stability gradient can be used to geographically illustrate the smallest change (measured by any suitably chosen norm) to a background flow that results in the largest change to a particular stability parameter (or linear combination of them). Thus, while the singular value decomposition seeks structures that grow most rapidly for a given background flow, the structural stability gradient indicates regions of the background flow where the stability properties of a particular perturbation structure are most sensitive. The computational effort of the structural stability gradient calculation is greatly reduced by an efficient means of estimating the first order

changes of the desired stability parameters. This subject will be investigated in a future study.

7 Acknowledgments

Discussions with our colleagues, especially Dr. C. Penland and Dr. J. Barsugli are gratefully acknowledged.

References

Borges, M. D. and P. D. Sardeshmukh, 1995. Barotropic Rossby wave dynamics of zonally varying upper-level flows during northern winter. *J. Atmos. Sci.*, **52**, 3779–3796.

Buizza, R. and T. N. Palmer, 1995: The singular-vector structure of the atmospheric global circulation. *J. Atmos. Sci.*, **52**, 1434–1456.

North, G. R., T. L. Bell, R. F. Cahalan, and F. J. Moeng, 1982. Sampling errors in the estimation of empirical orthogonal functions. *Mon. Wea. Rev.*, **110**, 699–706.

Penland, C. and P. D. Sardeshmukh, 1995. Error and sensitivity analysis of geophysical eigensystems. *J. Climate*, **8**, 1988–1998.

Trefethen, L. N., 1992: Pseudospectra of matrices. In: *Numerical Analysis 1991*. (ed. D. F. Griffiths and G. A. Watson). Longman, White Plains, NY, 234–266.

Trefethen, L. N., A. E. Trefethen, S. C. Reddy, and T. A. Driscoll, 1993. Hydrodynamic stability without eigenvalues. *Science*, **261**, 578–584.

Wilkinson, J. H., 1965. *The Algebraic Eigenvalue Problem*. Oxford University Press, 662.

List of Figures

Figure 1: Examples of the 250mb background flows used in this study, represented by streamfunction and isotachs. The upper panels show the (a) *NMC* 250mb and (b) *ECMWF* 250mb climatological flows. The lower panels show two states constructed by modifying the nonzonal component of the *NMC* climatology such that (c) α=0.8 and (d) α=1.2 in (23). Streamfunction is contoured in thin lines at intervals of 10^7 m^2s^{-1} throughout. Isotachs are drawn in thick lines every 15ms^{-1}, beginning at 30 ms^{-1}.

Figure 2: The relative change in the leading eigenvalue's (a) growth rate σ and (b) frequency ω as a function of the parameter γ that blends the *NMC* 250mb and *ECMWF* 250mb climatologies (see equation (22)). Solid bars correspond to the perturbation theory estimate, and hatched bars to the eigenvalues determined by the full matrix eigenanalysis.

Figure 3: (a) The fractional error $\|\hat{\underset{\sim}{a}} - \underset{\sim}{a}\|_\infty / \|\underset{\sim}{a}\|_\infty$ of the estimated eigenvector, where $\hat{\underset{\sim}{a}}$ is the vector of estimated expansion coefficients, $\underset{\sim}{a}$ is the vector of expansion coefficients obtained from the full matrix solution, and $\| \ \|_\infty$ is the infinity norm, $\|\underset{\sim}{\mathbf{x}}\|_\infty = max_i \ |x_i|$. Solid bars correspond to the **a**-phase and hatched bars to the **b**-phase of the eigenvector. (b) The maximum pattern correlation between the eigenvector estimated by perturbation theory and that obtained from a full matrix eigenanalysis.

Figure 4: The relative change in the leading singular value λ as a function of the parameter γ that blends the *NMC* 250mb and *ECMWF* 250mb climatologies. Solid bars correspond to the perturbation estimate and hatched bars to the full singular value decomposition.

Figure 5: (a) As in figure 3a, but for the estimated singular vector. (b) The maximum pattern correlation between the eigenvector estimated by perturbation theory and that obtained from the full matrix eigenanalysis.

Figure 6: The relative change in the leading eigenvalue's (a) growth rate σ and (b) frequency ω as a function of the parameter α that modulates the amplitude of the zonally asymmetric component of the *NMC* 250mb climatology (see equation (23)). Solid bars correspond to the perturbation theory estimate, and hatched bars to the eigenvalues determined by the full matrix eigenanalysis.

Figure 7: As in figure 3a, but as a function of the parameter α that modulates the amplitude of the zonally asymmetric component of the NMC 250mb climatology. (b) The maximum pattern correlation between the eigenvector estimated by perturbation theory and that obtained from a full matrix eigenanalysis.

Figure 8: The relative change in the leading singular value λ as a function of the parameter α that modulates the amplitude of the zonally asymmetric component of the NMC 250mb climatology. Shading as in figure 4.

Figure 9: (a) As in figure 5a, but as a function of the parameter α that modulates the amplitude of the zonally asymmetric component of the NMC 250mb climatology. (b) The pattern correlation between the leading singular vector estimated by perturbation theory and that obtained from a full matrix eigenanalysis as a function of α.

Figure 10: The average amplitude, $\bar{c}_l$, of the expansion coefficient matrix elements (see equation (14)) defined by

$$\bar{c}_l = \begin{cases} (1/(N-l))\sum_j |c_{j,\,j+l}|, & l \geq 0 \\ (1/(N+l))\sum_j |c_{j-l,\,j}|, & l < 0 \end{cases}$$

(*i.e.*, the average taken along each diagonal). The diagonal averages are shown for the leading singular vector coefficient matrices obtained in the (a) $\gamma = 0.95$ and (b) $\alpha = 1.3$ cases in equations (22) and (23), respectively.

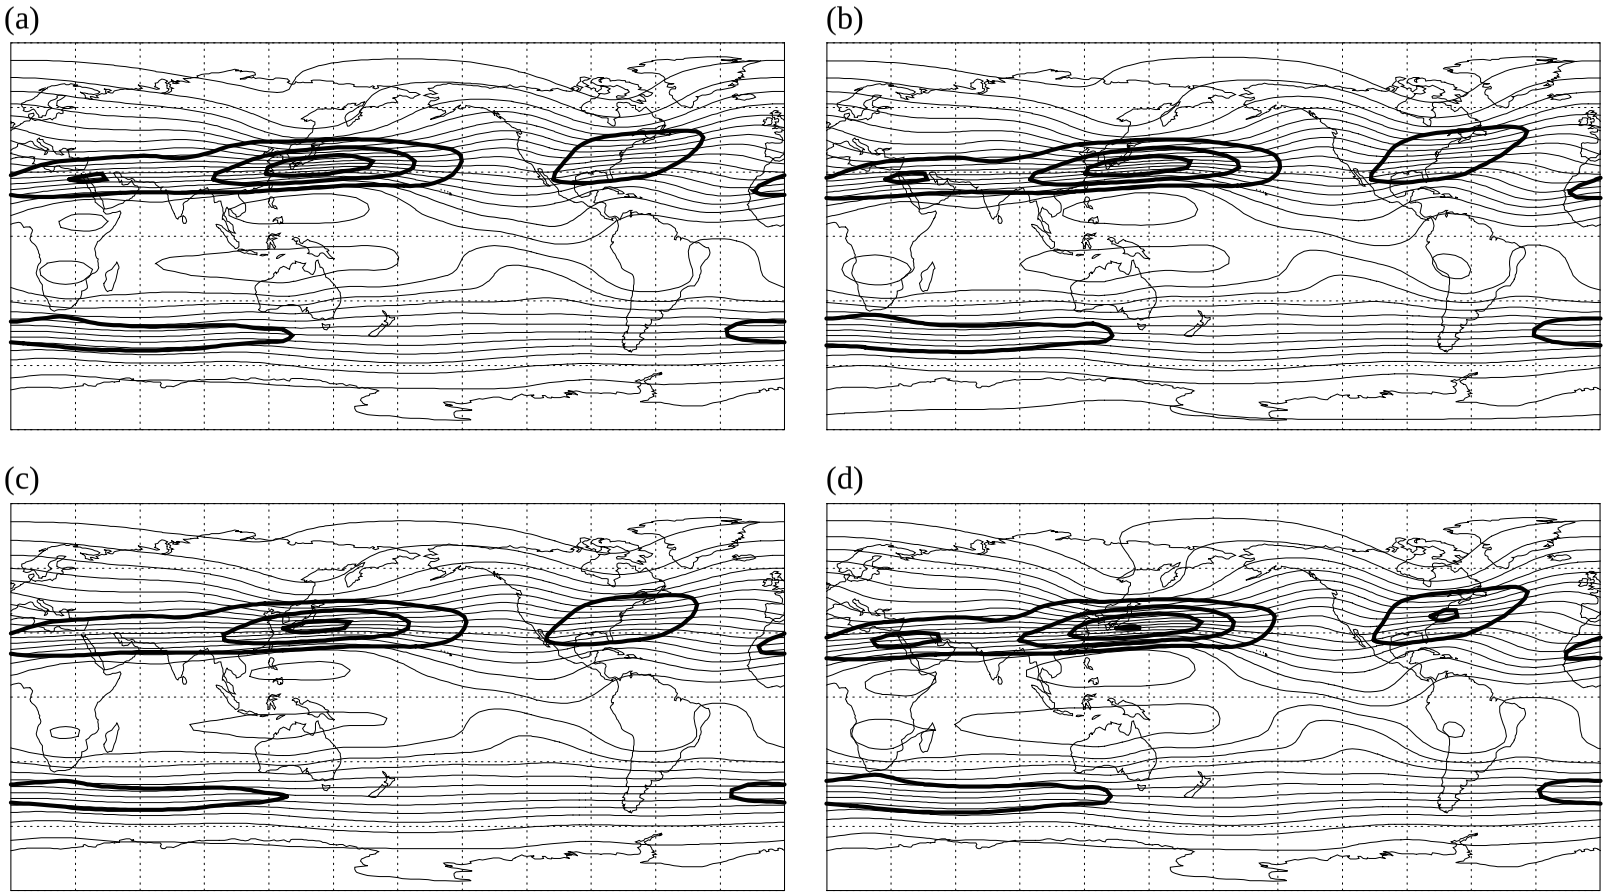

Figure 1: Examples of the 250mb background flows used in this study, represented by streamfunction and isotachs. The upper panels show the (a) *NMC* 250mb and (b) *ECMWF* 250mb climatological flows. The lower panels show two states constructed by modifying the nonzonal component of the *NMC* climatology such that (c) α=0.8 and (d) α=1.2 in (23). Streamfunction is contoured in thin lines at intervals of 10^7 m^2s^{-1} throughout. Isotachs are drawn in thick lines every 15ms^{-1}, beginning at 30 ms^{-1}.

$$\bar{\psi} = (1-\gamma)\bar{\psi}_{\text{NMC}} + \gamma\bar{\psi}_{\text{ECMWF}}, \qquad 0.05 \le \gamma \le 0.95$$

(a) Growth rate

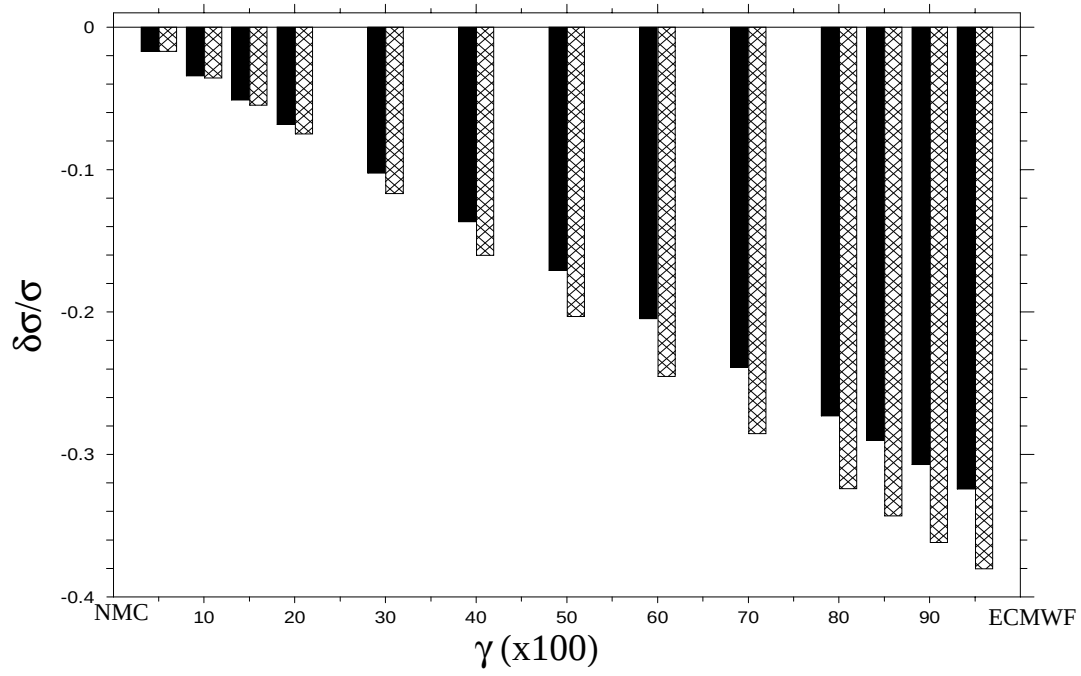

(b) Frequency

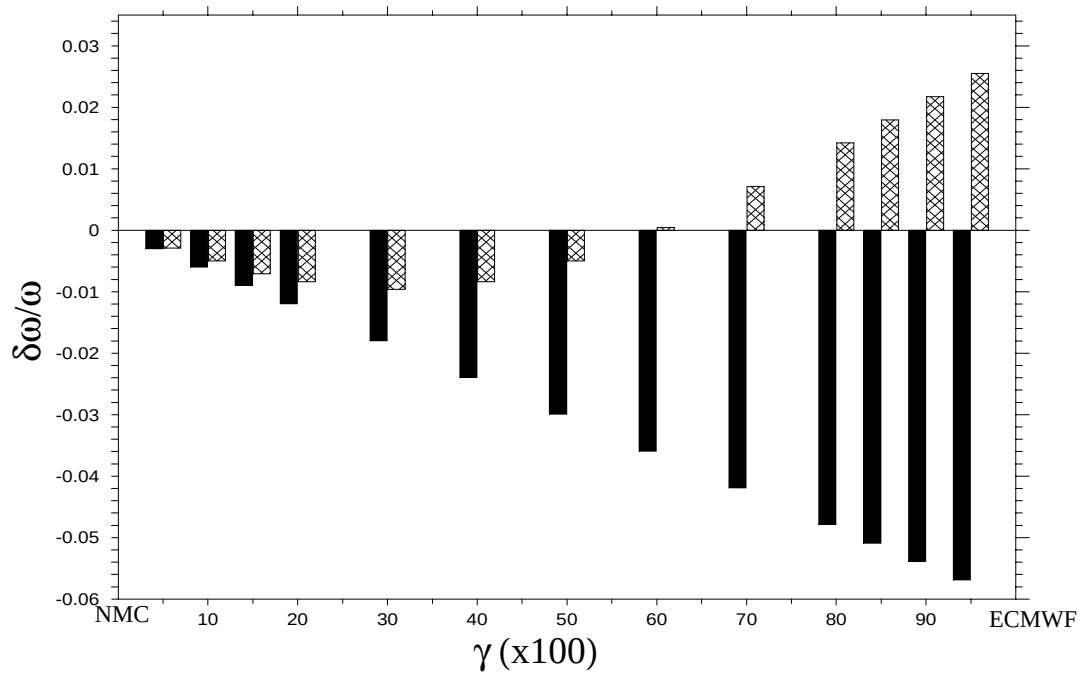

Figure 2: The relative change in the leading eigenvalue's (a) growth rate σ and (b) frequency ω as a function of the parameter γ that blends the NMC 250mb and ECMWF 250mb climatologies (see equation (22)). Solid bars correspond to the perturbation theory estimate, and hatched bars to the eigenvalues determined by the full matrix eigenanalysis.

$$\bar{\psi} = (1-\gamma)\bar{\psi}_{\mathrm{NMC}} + \gamma\bar{\psi}_{\mathrm{ECMWF}}, \qquad 0.05 \le \gamma \le 0.95$$

(a) Error norm

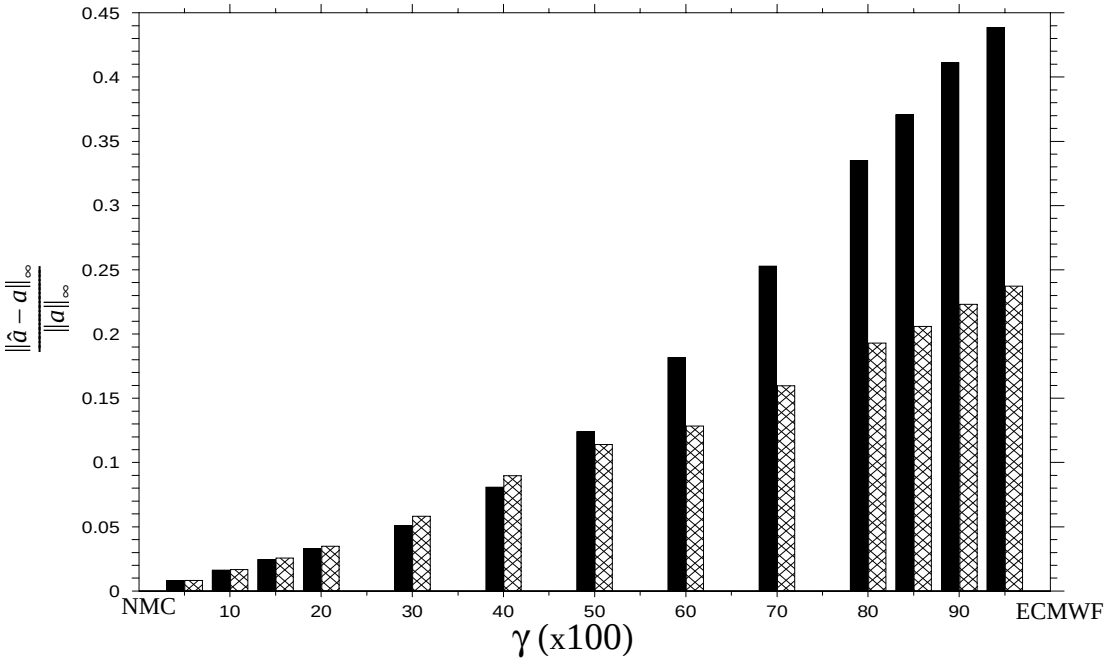

(b) Maximum pattern correlation

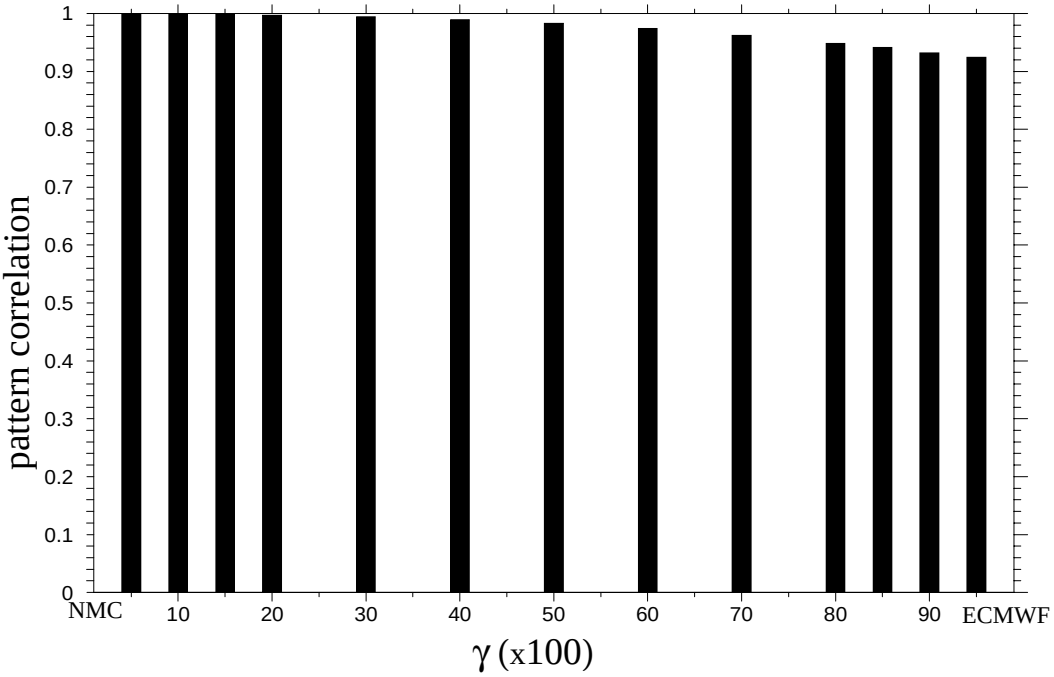

Figure 3: (a) The fractional error $\|\hat{a}-a\|_\infty / \|a\|_\infty$ of the estimated eigenvector, where $\hat{a}$ is the vector of estimated expansion coefficients, a is the vector of expansion coefficients obtained from the full matrix solution, and $\|\ \|_\infty$ is the infinity norm, $\|\mathbf{x}\|_\infty = max_i\ |x_i|$. Solid bars correspond to the **a**-phase and hatched bars to the **b**-phase of the eigenvector. (b) The maximum pattern correlation between the eigenvector estimated by perturbation theory and that obtained from a full matrix eigenanalysis.

$$\bar{\psi} = (1-\gamma)\bar{\psi}_{\mathrm{NMC}} + \gamma\bar{\psi}_{\mathrm{ECMWF}}, \qquad 0.05 \le \gamma \le 0.95$$

(a) Singular value

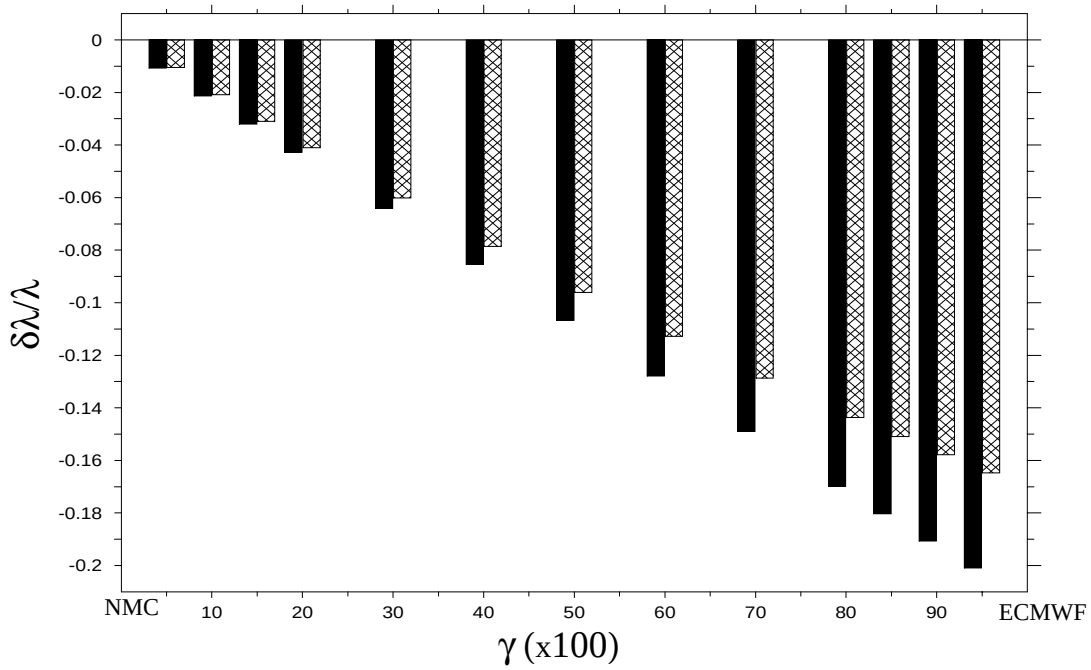

Figure 4: The relative change in the leading singular value λ as a function of the parameter γ that blends the NMC 250mb and ECMWF 250mb climatologies. Solid bars correspond to the perturbation estimate and hatched bars to the full singular value decomposition.

$$\bar{\psi} = (1-\gamma)\bar{\psi}_{\mathrm{NMC}} + \gamma\bar{\psi}_{\mathrm{ECMWF}}, \qquad 0.05 \le \gamma \le 0.95$$

(a) Error norm

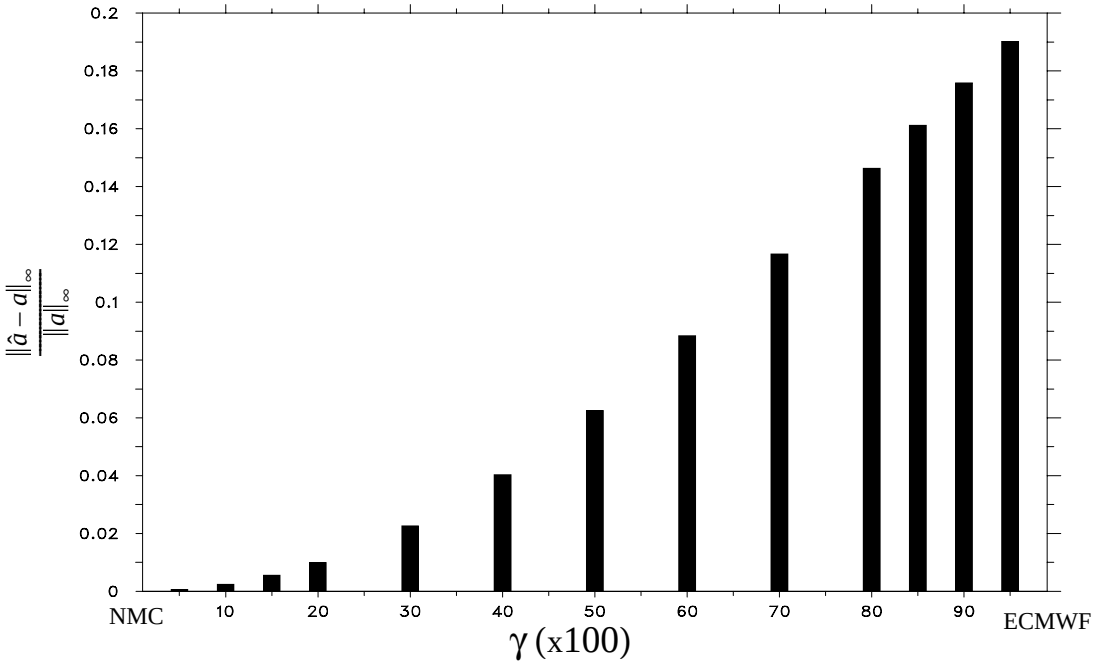

(b) Maximum pattern correlation

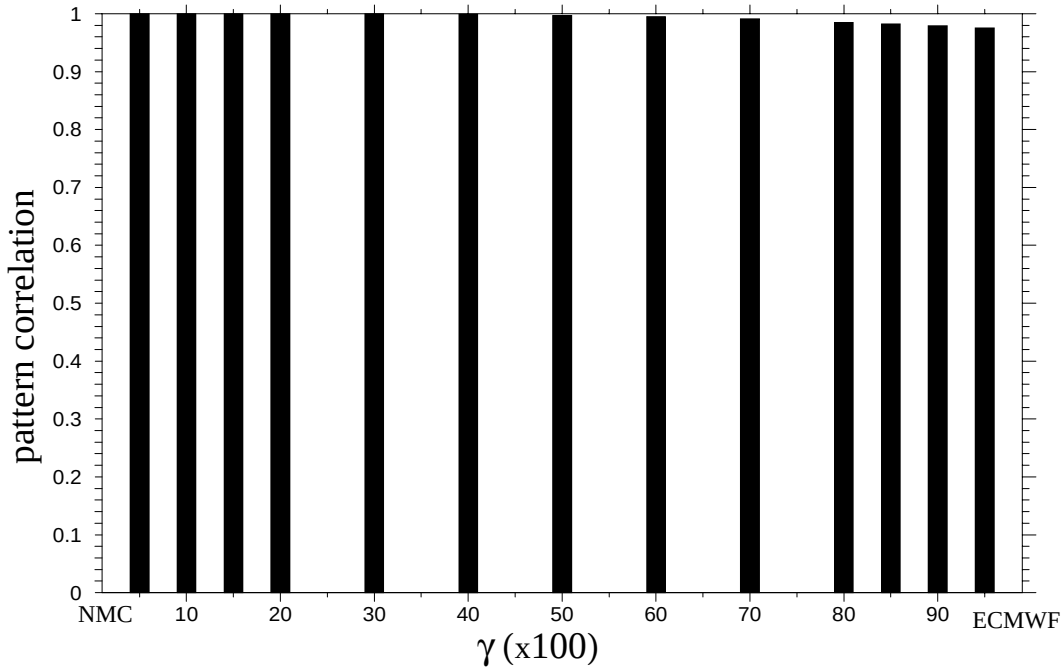

Figure 5: (a) As in figure 3a, but for the estimated singular vector. (b) The maximum pattern correlation between the eigenvector estimated by perturbation theory and that obtained from the full matrix eigenanalysis.

$$\bar{\psi} = [\bar{\psi}]_{NMC} + \alpha \bar{\psi}^*_{NMC} \quad , \quad 0.7 \le \alpha \le 1.6$$

(a) Growth rate

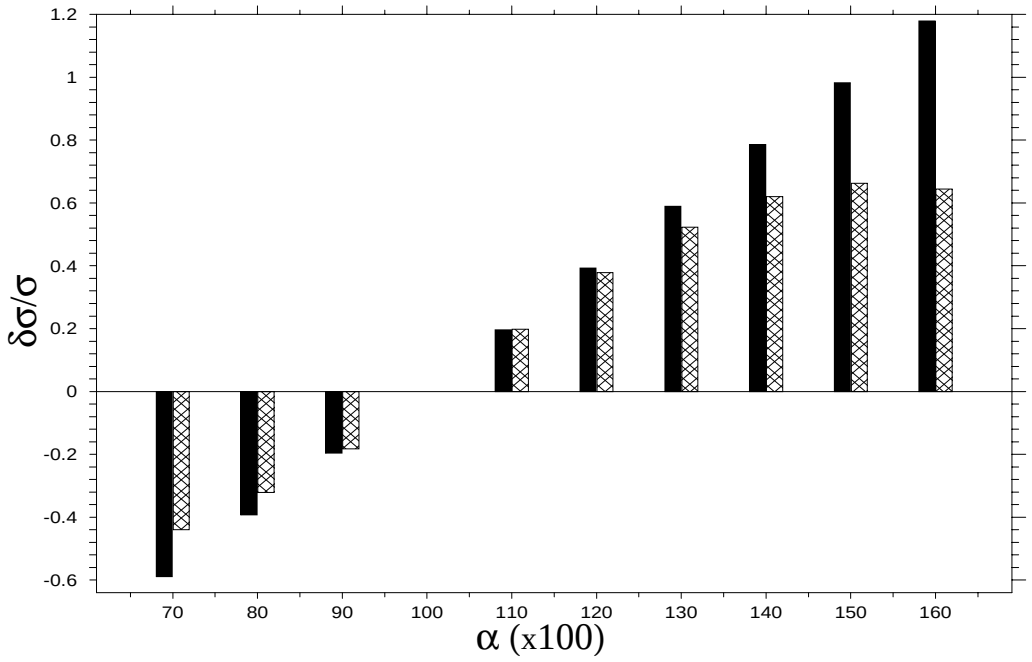

(b) Frequency

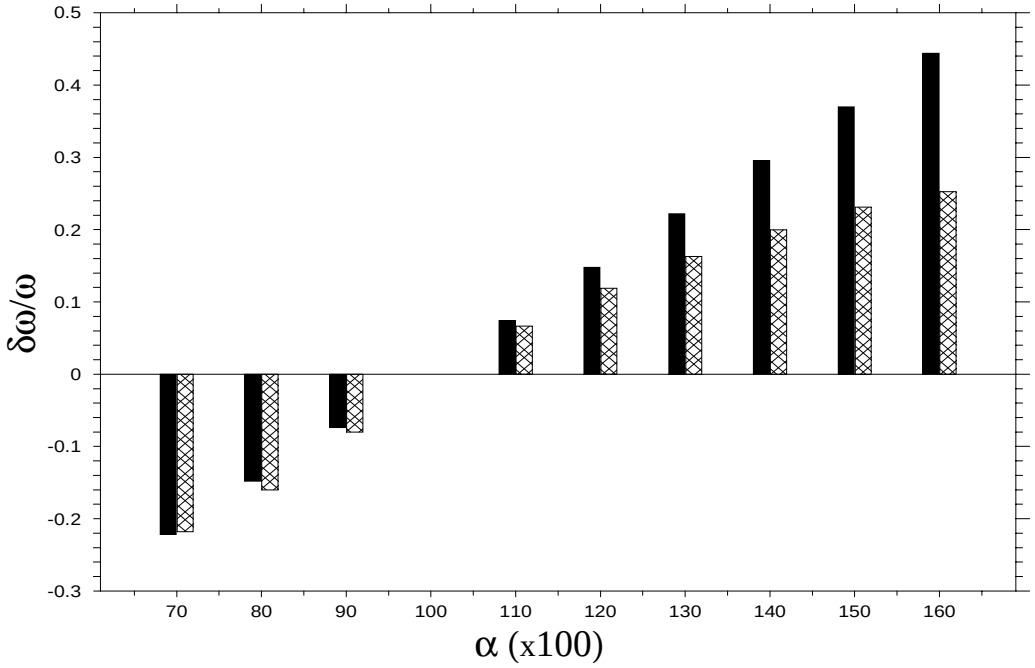

Figure 6: The relative change in the leading eigenvalue's (a) growth rate σ and (b) frequency ω as a function of the parameter α that modulates the amplitude of the zonally asymmetric component of the NMC 250mb climatology (see equation (23)). Solid bars correspond to the perturbation theory estimate, and hatched bars to the eigenvalues determined by the full matrix eigenanalysis.

$$\bar{\psi} = [\bar{\psi}]_{NMC} + \alpha\bar{\psi}^*_{NMC} \quad , \quad 0.7 \le \alpha \le 1.6$$

(a) Error norm

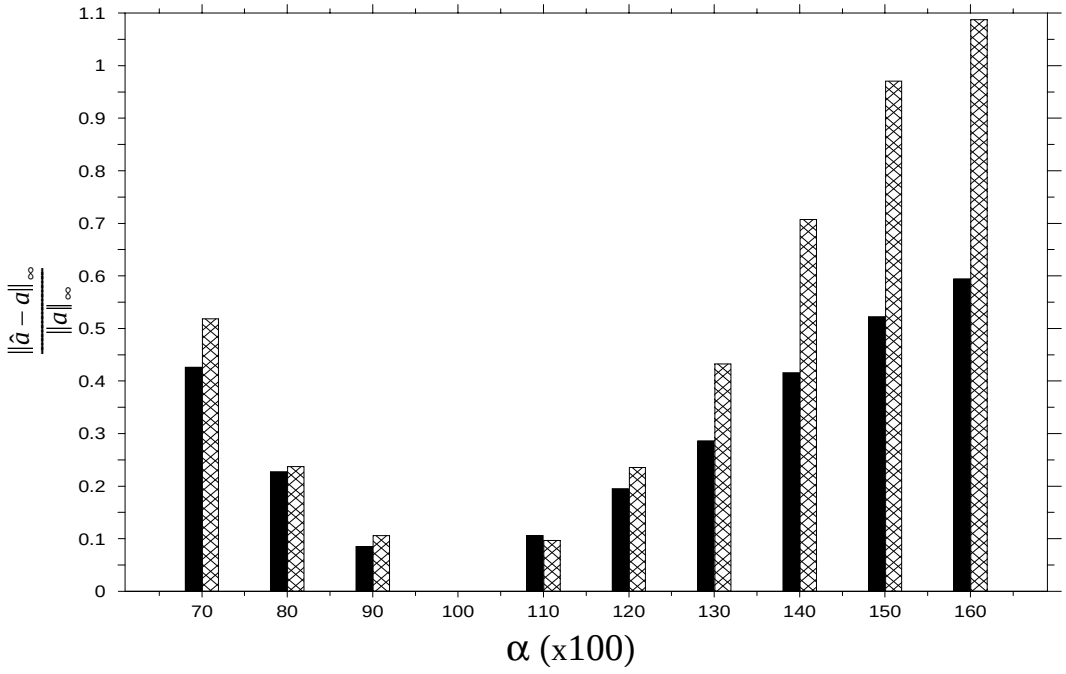

(b) Maximum pattern correlation

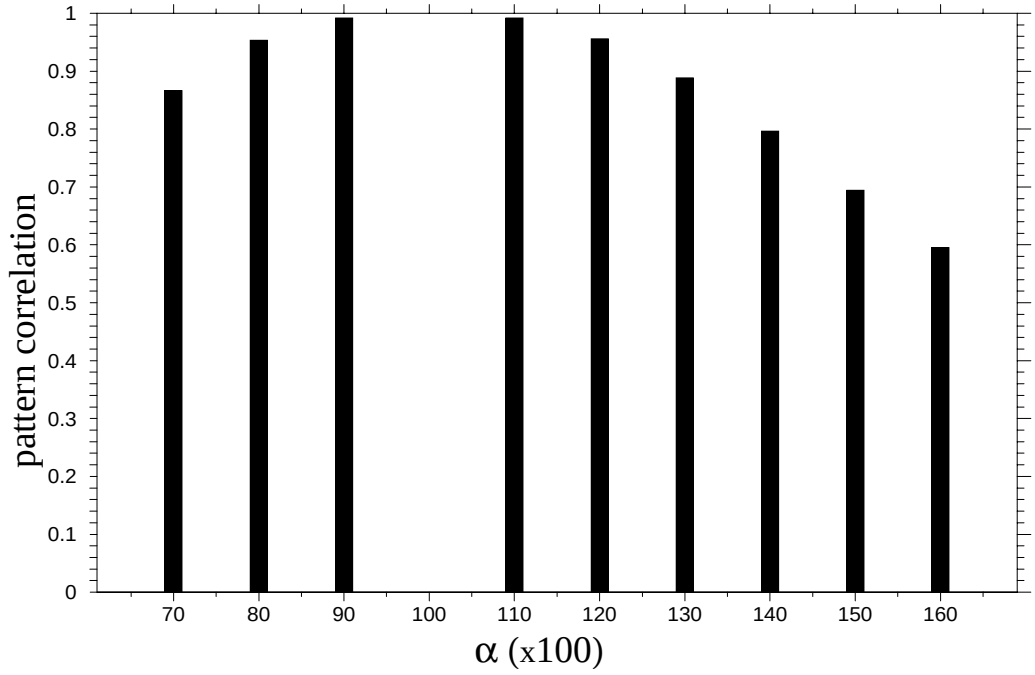

Figure 7: As in figure 3a, but as a function of the parameter α that modulates the amplitude of the zonally asymmetric component of the NMC 250mb climatology. (b) The maximum pattern correlation between the eigenvector estimated by perturbation theory and that obtained from a full matrix eigenanalysis.

$$\bar{\psi} = [\bar{\psi}]_{NMC} + \alpha \bar{\psi}^*_{NMC} \quad , \quad 0.7 \le \alpha \le 1.6$$

(a) Singular value

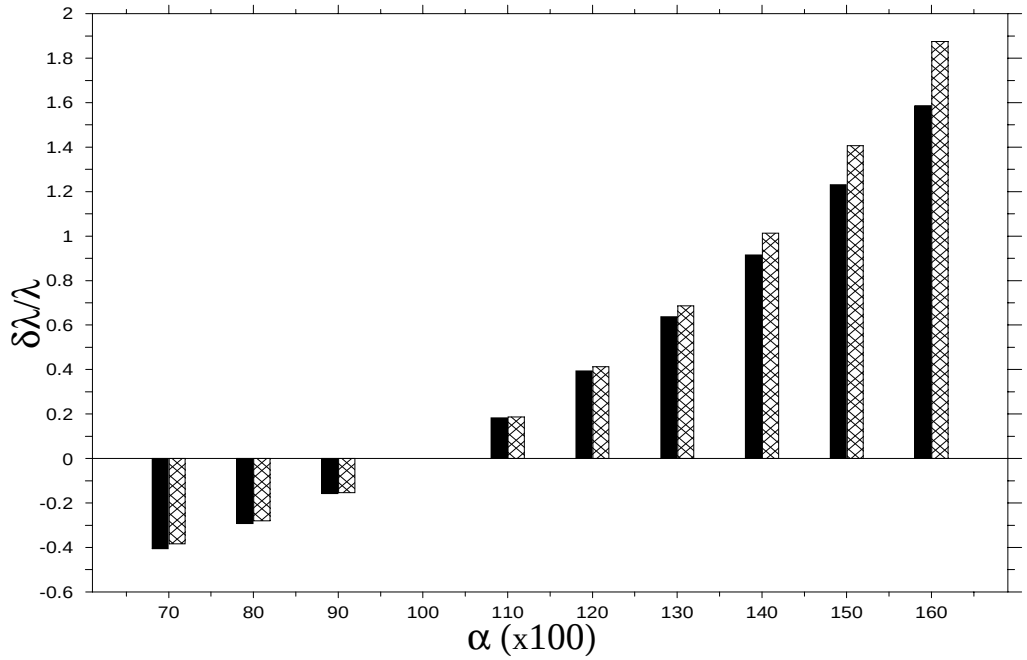

Figure 8: The relative change in the leading singular value λ as a function of the parameter α that modulates the amplitude of the zonally asymmetric component of the NMC 250mb climatology. Shading as in figure 4.

$$\bar{\psi} = [\bar{\psi}]_{NMC} + \alpha \bar{\psi}^*_{NMC} \quad , \quad 0.7 \le \alpha \le 1.6$$

(a) Error norm

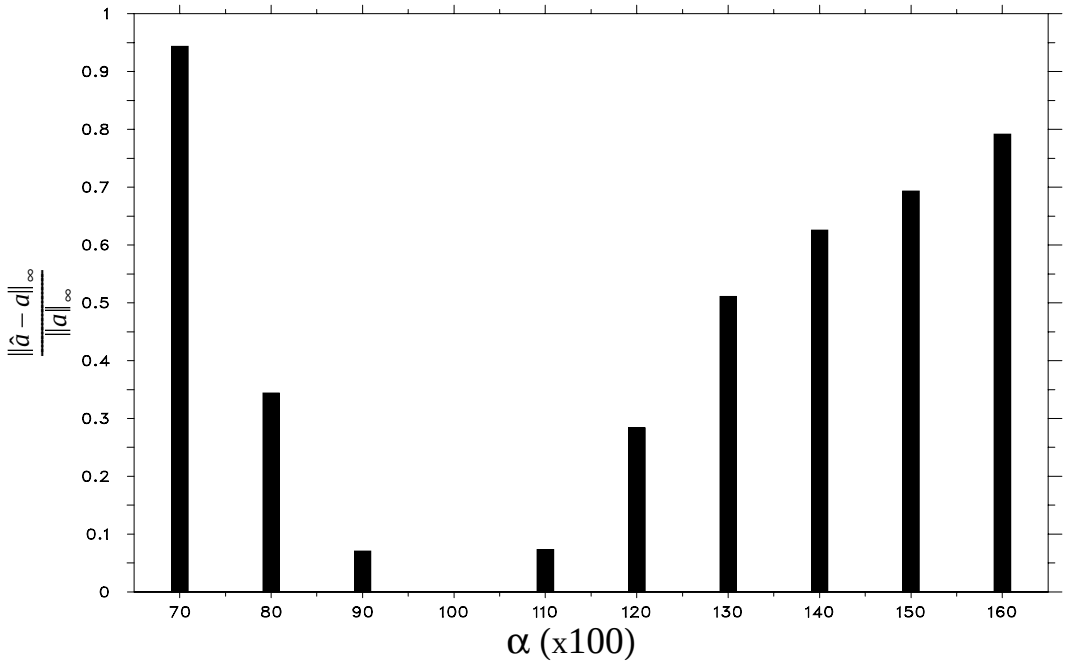

(b) Maximum pattern correlation

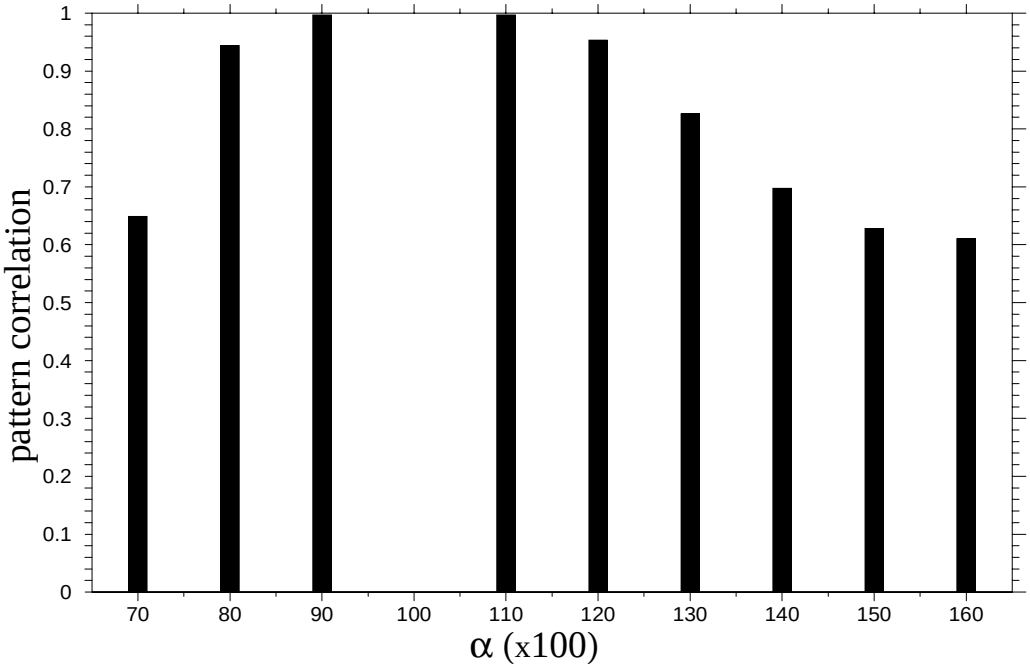

Figure 9: (a) As in figure 5a, but as a function of the parameter α that modulates the amplitude of the zonally asymmetric component of the NMC 250mb climatology. (b) The pattern correlation between the leading singular vector estimated by perturbation theory and that obtained from a full matrix eigenanalysis as a function of α.

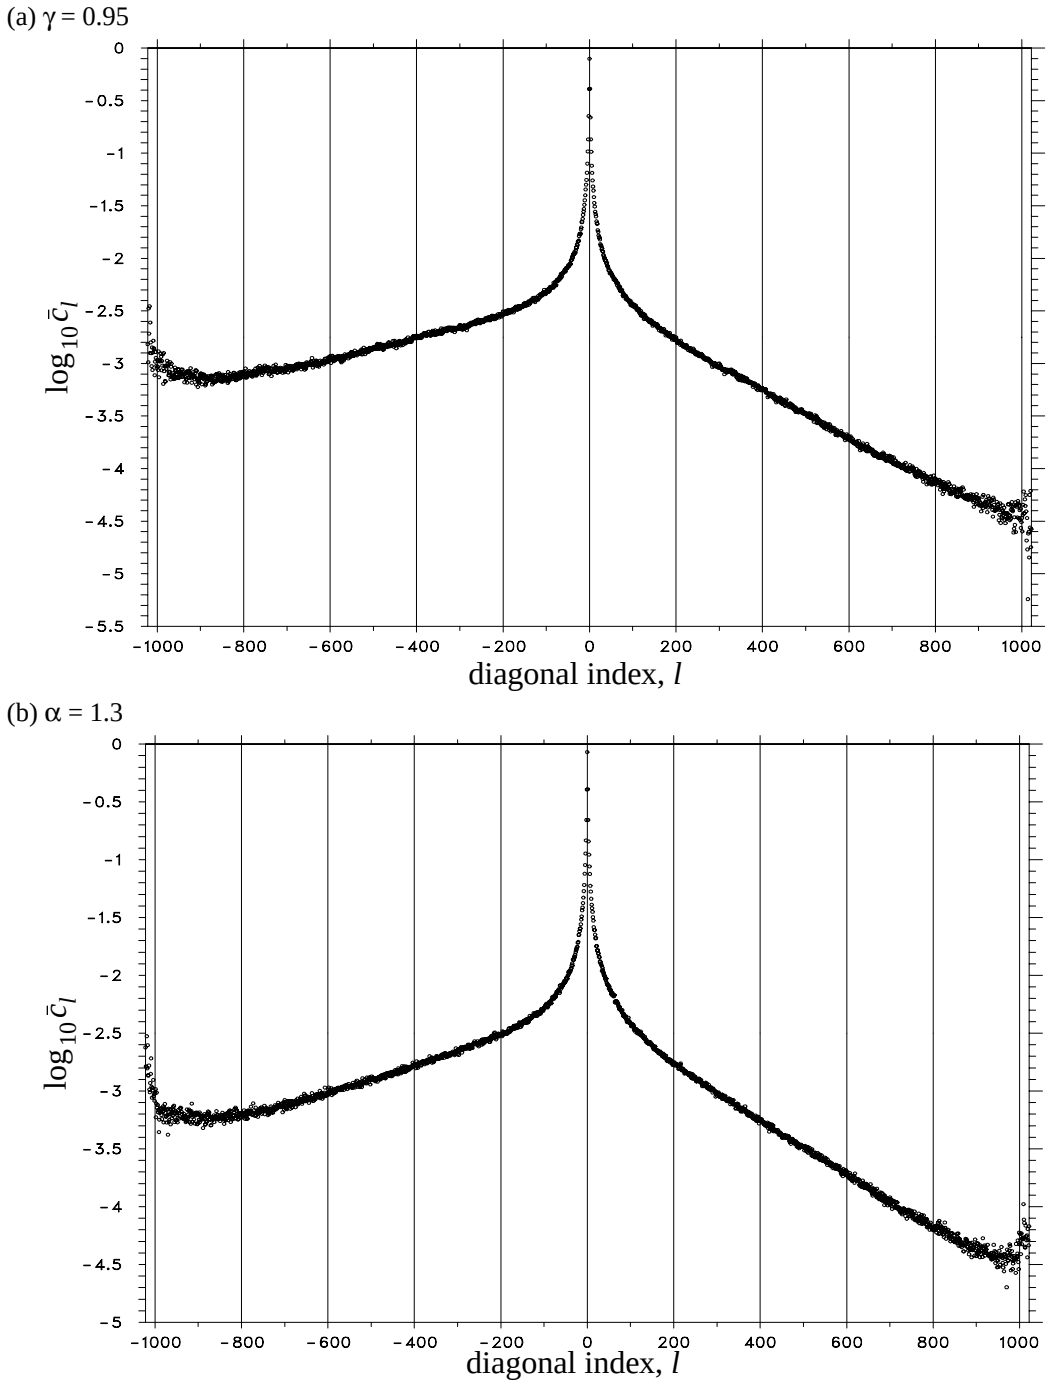

Figure 10: The average amplitude, $\bar{c}_l$, of the expansion coefficient matrix elements (see equation (14)) defined by

$$\bar{c}_l = \begin{cases} (1/(N-l))\sum_j \left| c_{j,\, j+l} \right|, & l \geq 0 \\ (1/(N+l))\sum_j \left| c_{j-l,\, j} \right|, & l < 0 \end{cases}$$

(*i.e.*, the average taken along each diagonal). The diagonal averages are shown for the leading singular vector coefficient matrices obtained in the (a) $\gamma = 0.95$ and (b) $\alpha = 1.3$ cases in equations (22) and (23), respectively.